Dynamic Scheduling for Charging Electric Vehicles: A Priority Rule

Yunjian Xu, Feng Pan, Lang Tong

arXiv:1602.00372v1 [math.OC] 1 Feb 2016

Abstract

We consider the scheduling of multiple tasks with pre-determined deadlines under random processing cost. This problem is motivated by the potential of large scale adoption of plug-in (hybrid) electric vehicles (PHEVs) in the near future. The charging requests of PHEVs usually have deadline constraints, and the electricity cost associated with PHEV charging is usually random due to the uncertainty in both system load and renewable generation. We seek to properly schedule the battery charging of multiple PHEVs so as to minimize the overall cost, which is derived from the total charging cost and the penalty for not completing charging before requested deadlines. Through a dynamic programming formulation, we establish the Less Laxity and Longer remaining Processing time (LLLP) principle that improves any charging policy on a sample-path basis, when the non-completion penalty is a convex function of the additional time needed to fulfill the uncompleted request. Specifically, the LLLP principle states that priority should be given to vehicles that have less laxity and longer remaining processing times. Numerical results demonstrate that heuristic policies that violate the LLLP principle, for example, the earliest deadline first (EDF) policy, can result in significant performance loss.

Index Terms

Plug-in hybrid electric vehicle, Dynamic programming, Demand response, Deadline scheduling

I. Introduction

We study the scheduling of multiple processors to perform tasks with deadlines under random processing cost. Each task requires a certain amount of processing time before its deadline, and failure to

Yunjian Xu is with the Engineering Systems and Design Pillar, Singapore University of Technology and Design, Singapore (e-mail: yunjian_xu@sutd.edu.sg). Feng Pan is with the Electricity Infrastructure Group, Pacific Northwest National Laboratory, Richland, WA 99354, USA (e-mail: feng.pan@pnnl.gov). Lang Tong is with the School of Electrical and Computer Engineering, Cornell University, Ithaca, NY 14853, USA (e-mail: ltong@cornell.edu).

This research was supported in part by the MIT-SUTD International Design Center (IDC) Grant IDG21400103, and the National Science Foundation under CNS 1135844. A preliminary version of this technical note appeared in the Proceedings of the 2012 IEEE Conference on Decision and Control [23].

fulfill its request incurs non-completion penalty. Different from most existing deadline scheduling models, our formulation allows the instantaneous processing cost to be time-variant and stochastic. We seek to characterize an optimal scheduling policy that minimizes the long-term expected total cost (the sum of processing cost and non-completion penalty). Although the results derived in this technical note generally apply to the aforementioned framework, we will focus on the scheduling of PHEV charging that may have significant impacts on both the reliability and efficiency of the next generation electric power grids.

Becoming popular in many countries, PHEVs (plug-in hybrid electric vehicles) may achieve significant market share over the next decade. However, the charging of a large number of PHEVs can add considerable stress to an existing power grid, especially at the distribution network level [7], [17]. The scheduling of charging PHEVs receives much attention in recent years [4], [22], [12], [13]. To minimize the load variance through PHEV charging, a few recent papers propose several approaches based on game theoretic analysis [18], [14] and decentralized optimization [11]. Although dynamic programming based approaches have been employed to study the optimal control of power management for a single PHEV [20], [19], there lacks a dynamic framework (on the scheduling of charging multiple PHEVs) that explicitly incorporates the stochasticity in both PHEV arrivals and charging costs.

This work is intimately related to the literature on deadline scheduling. For a single processor scheduling problem, it is well known that the earliest deadline first (EDF) policy [15] and the least-laxity first (LLF) policy [9] are optimal, if it is feasible to finish all tasks before their deadlines. When the completion of all tasks is not feasible, it has been demonstrated that EDF and LLF may perform poorly [16].[1] Closer to the present work, the authors of [3], [5] conduct a dynamic programming based approach to characterize optimal scheduling policies for the delivery of messages that would extinct after their individual deadlines. In the aforementioned literature processing capacity (of each individual processer) is usually assumed to be constant over the entire operation interval. As noted in [21], the scheduling of PHEV charging is fundamentally different, since the cost associated with PHEV charging is time-varying and stochastic (due to the inherent volatility in renewable generation and system load).

In this note, we consider a system with multiple (possibly a large number of) PHEVs and an underlying power grid with renewable generation. A system operator schedules the charging of PHEVs so as to minimize the long-run average cost. The formulated dynamic program (DP) incorporates *arbitrary* randomness in both the charging cost and the PHEV arrival processes.

The main contribution of this technical note is to establish an important and somewhat counter-intuitive

[1]There is also a substantial literature on deadline scheduling of multiple processors [10]; for a survey, see [8].

(partial) characterization on optimal scheduling policies. In particular, we show the **Less Laxity and Longer remaining Processing time** (LLLP) principle: priority should be given to vehicles that have less laxity and longer remaining processing times, if the non-completion penalty (as a function of the additional time needed to complete the task) is convex.[2] For a given heuristic policy, we show that an LLLP-based "interchanging" policy cannot be worse than the original heuristic. This result holds on every sample path and is robust against arbitrary random arrival process and charging cost. We also show (under some additional mild assumptions) the existence of an optimal stationary policy that always gives priority to vehicles with less laxity and longer remaining processing times. Numerical results presented in Section V show that the LLLP principle is practically useful: heuristic policies that violate the LLLP principle, such as the well known earliest deadline first (EDF) policy, can result in significant performance loss.

II. Model

We consider an infinite-horizon discrete time model. As in [21], we assume that each vehicle reports its arrival time, departure time, and charging request to the system operator at its arrival. The system operator uses all information available at the current stage (*i.e.*, the current system state of the DP to be formulated in Section III that includes the states of all arrived vehicles, the operating condition of the power system, as well as the prediction on future PHEV arrivals) to schedule the charging of PHEVs.

We study the scheduling problem of N PHEV chargers. For $i = 1, \ldots, N$, we refer to the vehicle that is connected to the ith charger as vehicle i. At stage t, let $\mathcal{I}_t \subseteq \{1, \ldots, N\}$ denote the set of chargers that are connected to electric vehicles, and $|\mathcal{I}_t|$ denote the number of vehicles connected to chargers. For each vehicle $i \in \mathcal{I}_t$, let α_i and β_i be its arrival and departure time, respectively. Under the assumption that both arrival and departure occur at the beginning of each stage, vehicle i can be charged from stage α_i through stage $\beta_i - 1$. We assume that $1 \leq \beta_i - \alpha_i \leq B$, i.e., every vehicle stays at a facility for at least one stage, and at most B stages.

For every $i \in \mathcal{I}_t$, let $\gamma_{i,t}$ denote its **remaining processing time** at stage t, i.e., the number of time units of charging needed to meet vehicle i's charging request under a time-invariant constant charging rate. We assume that the processing time of each vehicle is no greater than E. At stage t, for every $i \in \mathcal{I}_t$, we use a two-dimensional vector, $x_{i,t} \triangleq (\lambda_{i,t}, \gamma_{i,t})$, to denote the **state of vehicle** i, where $\lambda_{i,t} \triangleq \beta_i - t$

[2]According to the LLLP principle, for two vehicles with the same laxity, priority should be given to the vehicle with a later deadline (and longer remaining processing time). This is in sharp contrast to the case of a single processor with fixed processing capacity, where the earliest deadline first (EDF) policy is shown to be optimal.

is the number of remaining stages of the vehicle at charger i. For notational convenience, for $i \notin \mathcal{I}_t$ we let $x_{i,t} = (0, 0)$.

For every $i \in \mathcal{I}_t$, $a_{i,t} = 1$ if vehicle i is charged at stage t, and $a_{i,t} = 0$ otherwise. A feasible **action** at stage t, $\mathbf{a}_t = (a_{1,t}, \ldots, a_{N,t})$, is an N-dimensional vector with $a_{i,t} \leq \gamma_{i,t}$ for every i. Let A_t denote the total number of vehicles charged at stage t, *i.e.*, $A_t = \sum_{i \in \mathcal{I}_t} a_{i,t}$. For a vehicle $i \in \mathcal{I}_t$, if it remains connected to charger i at stage $t+1$, its state **evolves according to** $x_{i,t+1} = x_{i,t} - (1, a_{i,t})$. For an empty charger $i \notin \mathcal{I}_t$, if a vehicle arrives at stage $t+1$ then $i \in \mathcal{I}_{t+1}$ and the state of charger i becomes the initial state of this vehicle.

At stage t, let $s_t \in \mathcal{S}$ denote the **state of grid**, where the set $\mathcal{S}$ is assumed to be finite. The state of grid incorporates all the currently available information on all exogenous factors that have impacts on the cost associated with PHEV charging, such as the level of renewable generation and its prediction, the system load (excluding the PHEV charging load A_t) and its prediction, and the current time. The evolution of the state of grid depends on the current state s_t and the aggregate action[3] A_t. The **charging cost** at stage t $C(A_t, s_t)$ depends on the aggregate action A_t and the state of grid s_t.

At stage t, let $d_t \in \mathcal{D}$ be the **state of demand**, where $\mathcal{D}$ is assumed to be finite. The state of demand d_t contains all the currently available information on future PHEV arrivals, and completely determines the joint probability distribution on the number of arrival vehicles in the future and their initial states. The state of demand evolves as a time-homogeneous Markov chain, whose state transition is assumed to be independent of the state of the grid s_t and the action $\mathbf{a}_t$.

III. Dynamic Programming Formulation

In this section, we formulate the scheduling problem as an infinite-horizon dynamic program (DP) by introducing its state space, admissible action set, transition probabilities, stage cost, and average-cost objective function.

At each stage t, the **system state**, $\mathbf{x}_t$, consists of the states of all chargers, $\{x_{i,t}\}_{i=1}^N$, the state of grid s_t, and the state of demand d_t. Let $\mathcal{X}$ denote the set of all possible system states. Note that the size of state space grows exponentially with the number of chargers, N. Reasonable values of N, B, and E lead to very high dimensions, and make a direct solution to the DP impossible. We use $U_t(\mathbf{x}_t)$ to denote the set of feasible actions at stage t under system state $\mathbf{x}_t$.

[3]Note that charging a large number of vehicles may influence the Independent System Operator's (ISO) dispatch and reserve policy. To incorporate this type of impact, we allow the evolution of the grid state to depend on the aggregate action in general.

The **transition probability** of the system state depends on the current system state, $\mathbf{x}$, and the current action $\mathbf{a}_t$.[4] Since the state transition is independent of the stage index t, we use $p_{\mathbf{x},\mathbf{y}}(\mathbf{a}_t)$ to denote the transition probability from state $\mathbf{x}$ to $\mathbf{y}$, under the action $\mathbf{a}$.

At each stage t, the **stage cost** $g(\mathbf{x}_t, \mathbf{a}_t)$ consists of two parts: the charging cost $C(A_t, s_t)$ and the non-completion penalty. Let $\mathcal{J}(\mathbf{x}_t)$ denote the set of vehicles that will leave at stage $t+1$, i.e., $\mathcal{J}(\mathbf{x}_t) = \{j \in \mathcal{I}_t : \lambda_{j,t} = 1\}$. The stage cost function at stage t is[5]

$$g(\mathbf{x}_t, \mathbf{a}_t) = C(A_t, s_t) + \sum\nolimits_{j \in \mathcal{J}(\mathbf{x}_t)} q(\gamma_{j,t} - a_{j,t}), \tag{1}$$

where the penalty function $q : \mathbb{Z}_+ \to [0, \infty)$ with $q(0) = 0$ maps the number of uncharged battery units to its non-completion penalty (resulting from greenhouse gas emission or/and customers' inconvenience). Since both the set of system states and the set of feasible actions are finite, the stage cost is bounded.

A feasible **policy** $\pi = \{\nu_0, \nu_1, \dots, \}$ is a sequence of decision rules such that $\nu_t(\mathbf{x}_t) \in U_t(\mathbf{x}_t)$ for every t and $\mathbf{x}_t$. Given an initial system state $\mathbf{x}_0$, the time-averaged cost achieved by a policy π is given by

$$J_\pi(\mathbf{x}_0) \triangleq \limsup_{T \to \infty} \frac{1}{T} \mathbb{E}_\pi \left\{ \sum\nolimits_{t=0}^{T-1} g(\mathbf{x}_t, \nu_t(\mathbf{x}_t)) \right\}, \tag{2}$$

where the expectation is over the distribution of future system state $\{\mathbf{x}_t\}_{t=1}^{T-1}$ (induced by the policy π). Since the state evolution of the formulated DP does not depend on the time index and the state space is finite, there exists an optimal stationary policy $\pi^* = \{\mu^*, \mu^*, \dots\}$, and the limit on the right hand side of (2) exists [2].

Next we give an illustrative example of the general DP framework constructed above.

Example *3.1:* Our formulation incorporates the objective of minimizing load variance (that has been extensively explored in the literature [18], [11]). In this special case, the state of grid s_t is set to be the net system load (*i.e.*, the difference between system load and renewable generation) excluding PHEV charging. The charging cost is given by:

$$C(A_t, s_t) = H(A_t + s_t),$$

[4]Note that while the evolution of vehicles' states certainly depends on the action vector $\mathbf{a}_t$, the evolution of s_t depends only on the aggregate action A_t, and the evolution of d_t is completely exogenous.

[5]The state s_t can incorporate the maximum capacity constraint on all the N PHEV chargers by including an element c_t such that the charging cost becomes higher than the highest possible non-completion penalty $Nq(E)$ if $A_t > c_t$. On the other hand, our formulation omits the power flow constraints within a distribution network and cannot incorporate capacity constraints on any subset of the N chargers.

where $H(\cdot)$ is a strictly convex function that maps the total (net) system load to generation cost; a commonly used cost function is quadratic, e.g., $H(x) = x^2$ [18], [11]. Note that if the incremental non-completion penalty $q(n) - q(n-1)$ is set to be larger than the incremental charging cost $C(A_t, s_t) - C(A_t - 1, s_t)$, for every $n \geq 1$, $A_t \geq 1$, and $s_t \in \mathcal{S}$, the deadline requirement of each vehicle becomes a "hard constraint", in that it is optimal to fulfill all charging requests before their deadlines, as long as it is feasible to do so. ■

Remark *3.1:* Although the state of grid s_t and the state of demand d_t are modeled as stationary Markov chains, it is worth noting that the time dependency of the grid status (e.g., renewable generation and system load) and PHEV arrivals can be incorporated by including in the states s_t and d_t a periodic Markov chain that describes the evolution of local time.

The DP framework constructed in Sections II and III is general. The only conditions on s_t required by the LLLP principle (that will be formally stated and proved in Theorem 4.1) are: i) the charging cost at each stage t is of the form $C(A_t, s_t)$, which depends only on the aggregate action A_t and s_t, and ii) the evolution of s_t depends only on A_t (but not on $a_{i,t}$ for any i). In other words, the LLLP principle holds regardless of the detailed model used by the operator to describe the power grid dynamics (e.g., information included in the state s_t, its evolution, and the exact form of charging cost). ■

IV. The LLLP Principle

In this section we establish the main result of this technical note. In Section IV-A, we first define a partial order over the set of vehicle states: a vehicle with less laxity and a longer remaining processing time has a higher-order state. For any given (possibly non-stationary) heuristic policy that violates the LLLP principle, we construct an interchanging policy that gives priority to the vehicle with a higher-order state. We show that on every sample path, the interchanging policy can only reduce the ex-post (realized) cost, compared with the original heuristic (cf. Theorem 4.1). In Section IV-B, under some mild assumptions on the evolutions of the grid state s_t and the state of demand d_t, we show the existence of an optimal stationary policy that follows the LLLP principle.

A. LLLP-based Interchanging Policy

For every vehicle $i \in \mathcal{I}_t$, its **laxity** (at stage t) is defined by

$$\theta_{i,t} = \begin{cases} \lambda_{i,t} - \gamma_{i,t}, & \text{if } \gamma_{i,t} > 0, \\ B, & \text{if } \gamma_{i,t} = 0. \end{cases} \tag{3}$$

Note that for a vehicle i with $\gamma_{i,t} > 0$, its laxity $\theta_{i,t} \in \{1-E, 2-E, \ldots, B-1\}$ is the maximum number of stages it can tolerate before the time it has to be put on uninterrupted battery charging. We are now ready to define a partial order over the set of all possible vehicle states.

Definition *4.1:* For two vehicles $i, j \in \mathcal{I}_t$, we say $i \preccurlyeq j$ (vehicle j has priority over i) if j has less laxity and longer remaining processing time, i.e., $\theta_{i,t} \geq \theta_{j,t}$, $\gamma_{i,t} \leq \gamma_{j,t}$, and at least one of these two inequalities strictly holds. ■

It is not hard to check that the relation $\preccurlyeq$ is reflexive, antisymmetric, and transitive, and therefore is a partial order. We also note that if vehicle j has less laxity and a later deadline than vehicle i, i.e., if $\theta_{i,t} \geq \theta_{j,t}$ and $\lambda_{i,t} \leq \lambda_{j,t}$, then we must have $i \preccurlyeq j$.

At a system state $\mathbf{x}_t$, two vehicles i and j are *incomparable*, if $\theta_{i,t} \geq \theta_{j,t}$ and $\gamma_{i,t} > \gamma_{j,t}$, or $\theta_{i,t} > \theta_{j,t}$ and $\gamma_{i,t} \geq \gamma_{j,t}$. In this case, which vehicle should have higher priority depends on future system dynamics. On the other hand, if vehicle j has priority over vehicle i, we argue that priority should always be given to vehicle j, regardless of future system dynamics. This result requires the penalty function to be convex, as stated in the following assumption.

Assumption *4.1:* The incremental non-completion penalty is non-negative and non-decreasing, i.e.,

$$0 \leq q(n) - q(n-1) \leq q(n+1) - q(n), \quad n = 1, 2, \ldots.$$

The non-completion penalty may come from the inconvenience caused to customers as well as the potential environmental damage caused by the emission of PHEVs' combustion engines. We note that environmental damage is usually considered to be convex with respect to greenhouse gas emission [6].

Definition *4.2 (An LLLP-based Interchanging Policy):* Suppose that at some system state $\mathbf{x}_t$, vehicle j has priority over i, and that a policy $\pi = \{\nu_0, \nu_1, \ldots\}$ charges vehicle i but not j. Let $W \stackrel{\Delta}{=} \max\{\lambda_{i,t}, \lambda_{j,t}\} - 1$. We now formally define the **interchanging policy** $\bar{\pi} = \{\nu_0, \ldots, \nu_{t-1}, \bar{\nu}_t, \bar{\nu}_{t+1}, \ldots\}$ (generated from the policy π with respect to vehicles i and j at state $\mathbf{x}_t$) as follows.

1) We first let $\bar{\nu}_k = \nu_k$ for $k \geq t$, and then update the sequence of decision rules $\{\bar{\nu}_k\}_{k=t}^{t+W}$ as follows.
2) Policy $\bar{\pi}$ charges j instead of i at state $\mathbf{x}_t$. That is, $\bar{\nu}_t(\mathbf{x}_t)$ is the same as $\nu_t(\mathbf{x}_t)$ except that its ith component is 0 and its jth component is 1.

   Following the state $\mathbf{x}_t$, for any (realized) sequence of system states that would occur with positive probability under the policy π, $\{\mathbf{x}_k\}_{k=t+1}^{t+W}$, there exists a corresponding sequence of system states following the state-action pair $(\mathbf{x}_t, \bar{\nu}_t(\mathbf{x}_t))$, $\{\hat{\mathbf{x}}_k\}_{k=t+1}^{t+W}$. The corresponding state $\hat{\mathbf{x}}_k$ differs from $\mathbf{x}_k$ only on the states of vehicles i and j: $\hat{\gamma}_{i,k} = \gamma_{i,k} + 1$ for $k = t+1, \ldots, \beta_i - 1$, and $\hat{\gamma}_{j,k} = \gamma_{j,k} - 1$ for $k = t+1, \ldots, \beta_j - 1$.

3) For every $\{\mathbf{x}_k\}_{k=t+1}^{t+W}$, let $\mathcal{G}(\{\mathbf{x}_k\}_{k=t+1}^{t+W}) \subseteq \{t+1,\ldots,\min\{\beta_i,\beta_j\}-1\}$ be the set of stages that policy π charges vehicle j but not i, before vehicle i's departure. If the set $\mathcal{G}(\{\mathbf{x}_k\}_{k=t+1}^{t+W})$ is empty, let[6] $\bar{\nu}_k(\hat{\mathbf{x}}_k) = \nu_k(\mathbf{x}_k)$, for $k = t+1,\ldots,t+W$, i.e., the interchanging policy $\bar{\pi}$ agrees with the original policy π after stage t.

   If $\mathcal{G}(\{\mathbf{x}_k\}_{k=t+1}^{t+W})$ is not empty, let w be its minimal element. At stages $k = t+1,\ldots,w-1$, let $\bar{\nu}_k(\hat{\mathbf{x}}_k) = \nu_k(\mathbf{x}_k)$. At stage w, policy $\bar{\pi}$ charges vehicle i instead of j, i.e., $\bar{\nu}_w(\hat{\mathbf{x}}_w)$ is the same as $\nu_w(\mathbf{x}_w)$ except that its ith component is 1 and its jth component is 0. ■

Lemma *4.1:* An interchanging policy $\bar{\pi}$ is feasible.

The proof of Lemma 4.1 is given in Appendix A, where we show that the action taken by policy $\bar{\pi}$ at every system state is feasible.

Theorem *4.1 (The LLLP Principle):* Suppose that Assumption 4.1 holds and that, at some system state $\mathbf{x}_t$, vehicle j has priority over i in the sense of Definition 4.1. Let $\bar{\pi}$ be the interchanging policy generated from a policy π according to Definition 4.2. For every $T \geq \max\{\lambda_{j,t},\lambda_{i,t}\}-1$ and along every sample path from stage $t+1$ through stage $t+T$, the total (realized) cost resulting from the interchanging policy $\bar{\pi}$ cannot be higher than that achieved by the original policy π. □

Proof: It follows from Definition 4.2 that the interchanging policy always charges an equal number of vehicles as the original policy. Since the evolution of $\{s_t\}$ depends only on the total number of charged vehicles at each stage t, the two policies result in the same system dynamics. As a result, following the state $\mathbf{x}_t$, for every sequence of system states that would occur with positive probability under the policy π, $\{\mathbf{x}_k\}_{k=t+1}^{t+W}$, there exists a corresponding sequence of system states, $\{\bar{\mathbf{x}}_k\}_{k=t+1}^{t+W}$, which will occur with equal probability under policy $\bar{\pi}$. If the set[7] $\mathcal{G}(\{\mathbf{x}_k\}_{k=t+1}^{t+W})$ is empty, then $\{\bar{\mathbf{x}}_k\}_{k=t+1}^{t+W} = \{\hat{\mathbf{x}}_k\}_{k=t+1}^{t+W}$; otherwise, we have

$$\{\bar{\mathbf{x}}_k\}_{k=t+1}^{t+W} = \{\hat{\mathbf{x}}_{t+1},\ldots,\hat{\mathbf{x}}_w,\mathbf{x}_{w+1},\ldots,\mathbf{x}_{t+W}\},$$

where w be the minimum element in the set $\mathcal{G}(\{\mathbf{x}_k\}_{k=t+1}^{t+W})$.

We further note that the two policies, π and $\bar{\pi}$, are identical after stage $t+W$. As a result, to prove this theorem, it suffices to show that for every realization of system states under the policy π, $\{\mathbf{x}_k\}_{k=t+1}^{t+W}$,

[6]Lemma 4.1 shows that for $k = t+1,\ldots,t+W$, whenever the policy π charges vehicle j at state $\mathbf{x}_k$, it is feasible to charge vehicle j at the corresponding state $\hat{\mathbf{x}}_k$, i.e., $\hat{\gamma}_{j,k} \geq 1$.

[7]The set $\mathcal{G}(\{\mathbf{x}_k\}_{k=t+1}^{t+W})$ and the sequence $\{\hat{\mathbf{x}}_k\}_{k=t+1}^{t+W}$ are introduced in Definition 4.2.

and the corresponding realization of system states under the policy $\bar{\pi}$, $\{\bar{\mathbf{x}}_k\}_{k=t+1}^{t+W}$,

$$g(\mathbf{x}_t, \bar{\nu}_t(\mathbf{x}_t)) + \sum_{k=t+1}^{t+W} g(\bar{\mathbf{x}}_k, \bar{\nu}_k(\bar{\mathbf{x}}_k)) \leq g(\mathbf{x}_t, \nu_t(\mathbf{x}_t)) + \sum_{k=t+1}^{t+W} g(\mathbf{x}_k, \nu_k(\mathbf{x}_k)). \tag{4}$$

We now prove Eq. (4) by discussing the following two cases:

1. If $\mathcal{G}(\{\mathbf{x}_k\}_{k=t+1}^{t+W})$ is not empty, for every pair of system state realizations, $\{\mathbf{x}_k\}_{k=t+1}^{t+W}$ and $\{\bar{\mathbf{x}}_k\}_{k=t+1}^{t+W}$, both polices must result in the same ex-post cost, i.e., the equality holds in (4).

2. If $\mathcal{G}(\{\mathbf{x}_k\}_{k=t+1}^{t+W})$ is empty, then whenever π charges j, it must also charge i, for $k = t+1, \ldots, \min\{\beta_i, \beta_j\}-1$. For a sequence of system states realized under π, $\{\mathbf{x}_k\}_{k=t+1}^{t+W}$, let ρ_i denote the remaining processing time of vehicle i at its deadline β_i. That is, $\rho_i \triangleq \gamma_{i,\beta_i-1} - a_{i,\beta_i-1}$, where a_{i,β_i-1} is the action on vehicle i at state $\mathbf{x}_{\beta_i-1}$ according to policy π. Similarly, $\bar{\rho}_i$ is defined for the corresponding sequence $\{\bar{\mathbf{x}}_k\}_{k=t+1}^{t+W}$ under the interchanging policy $\bar{\pi}$. Since j has priority over i at $\mathbf{x}_t$, it is straightforward to check (from the definition of $\bar{\pi}$ in Definition 4.2) that $0 \leq \rho_i < \rho_j$, $\bar{\rho}_j = \rho_j - 1$, and $\bar{\rho}_i = \rho_i + 1$. It follows from Assumption 4.1 that $q(\bar{\rho}_j) + q(\bar{\rho}_i) \leq q(\rho_j) + q(\rho_i)$. Note that this inequality implies the desired result in (4), since the two policies, π and $\bar{\pi}$, result in the same cost except possible different penalties for not fulfilling vehicle i's and j's charging requests. ■

Before ending this section we make some brief discussion on the intuition behind the LLLP principle. We consider a simple **two-vehicle example**. At stage 0, the states of the two vehicles are $x_{1,0} = (2, 1)$ and $x_{2,0} = (3, 2)$. Note that vehicle 2 has priority to vehicle 1, according to the LLLP principle. If vehicle 1 is charged at stage 0, then the vehicle is fully charged (and not available for charging) at stage 1; on the other hand, if only vehicle 2 is charged at stage 0, then both vehicles are available for charging at stage 1. The LLLP principle argues that the latter situation is preferable, because

- in the latter situation, the operator has a larger set of feasible actions at stage 1;
- under convex penalty functions, remaining processing time should be split among multiple vehicles.

Under random charging cost and convex non-completion penalty, it is always desirable to have a larger number of smaller unfinished tasks that can be processed simultaneously when charging cost becomes lower in the future.

Remark *4.1:* Although an interchanging policy $\bar{\pi}$ cannot be worse than the original heuristic, it may still be (sometimes obviously) suboptimal, since it does not fully utilize the extra "flexibility" provided by the LLLP principle. Note that $\bar{\pi}$ charges i but not j at stage w (cf. Definition 4.2). A natural way to improve $\bar{\pi}$ is to charge both i and j at stage w under certain circumstances, e.g., when the laxity of vehicle j is small or it is cheap to charge an additional vehicle at stage w.

This intuition can be illustrated by the aforementioned two-vehicle example. The charging cost at stage 0, 1, 2 is A_0, 0, and $2A_2$, respectively (here A_t denotes the number of vehicles charged at stage t). Consider an EDF policy π that charges vehicle 1 at stage 0, and charges vehicle 2 at stages 1 and 2. According to Definition 4.2, $t = 0$ (when policy π violates the LLLP principle) and $w = 1$. The unique optimal policy gives priority to vehicle 2 at stage t (following the LLLP principle), and charges both vehicles at stage w. ■

B. Optimality of the LLLP Principle

In this subsection, we show the existence of an optimal stationary policy that always follows the LLLP principle. The following technical assumption is made to guarantee that the minimum average cost does not depend on the initial state.

Assumption *4.2:* We assume the following.

4.2.1. Under every $d \in \mathcal{D}$, there is positive probability that no vehicle arrives at the next stage.

4.2.2. There exists a special state of grid $\bar{s} \in \mathcal{S}$ such that for some positive integer m, for every initial state $s_0 \in \mathcal{S}$, and under the sequence of zero aggregated charging decisions, $\{A_0 = 0, A_1 = 0, \ldots, A_{m-1} = 0\}$, state $\bar{s}$ is visited with positive probability at least once within the first m steps.

4.2.3. The state of demand $\{d_t\}$ evolves as an ergodic Markov chain.

Note that Assumption 4.2.2 holds if without PHEV charging, (e.g., in a model with $N = 0$), the state of grid evolves as an ergodic Markov chain. Although Assumption 4.2.3 requires that $\{d_t\}$ is ergodic, the time dependency of PHEV arrival process can be incorporated by including in the state d_t a periodic Markov chain that describes time evolution.

Lemma *4.2:* Suppose that Assumption 4.2 holds. The minimum average cost is equal for all initial states, i.e.,

$$\lambda \stackrel{\Delta}{=} J_{\mu^*}(\mathbf{x}), \quad \forall\, \mathbf{x} \in \mathcal{X}, \tag{5}$$

where $J_{\mu^*}(\mathbf{x})$ is the (minimum) average cost achieved by an optimal stationary policy μ^* (cf. its definition in Eq. (2)).

Proof: We pick up a state $\bar{d} \in \mathcal{D}$, and define a **special system state**

$$\bar{\mathbf{x}} = \{(0,0), \ldots, (0,0), \bar{s}, \bar{d}\}, \tag{6}$$

where $\bar{s} \in \mathcal{S}$ is the special state of grid defined in Assumption 4.2.2. At this special system state, all charging facilities are empty with state $(0,0)$. According to Proposition 7.4.1 of [1], to show the desired

result we only need to argue that, for every initial system state $\mathbf{x}_0 \in \mathcal{X}$ and under all policies, the special system state is visited with positive probability at least once within the first $L \stackrel{\Delta}{=} B + \max\{m, |\mathcal{D}|\}$ steps, with m being the integer defined in Assumption 4.2.2.[8]

Under Assumption 4.2.1, the probability that no vehicle arrives within the first $L+1$ stages is positive, and if this is the case, all charging facilities are empty from stage B to stage L, regardless of the policy used by the operator. In this case, since no vehicle is charged from stage B to stage $L-1$, Assumption 4.2.2 implies that with positive probability the special state $\bar{s}$ is visited at least once from stage B to stage L. Also, since $\{d_t\}$ is an ergodic Markov chain, the probability that the state $\bar{d}$ is visited at least once from stage B to stage L is positive. Since the evolutions of $\{d_t\}$ and $\{s_t\}$ are assumed to be independent, for all initial system states and under all policies, the special system state $\bar{\mathbf{x}}$ is visited with positive probability at least once from stage B to stage L. ■

Theorem *4.2:* Suppose that Assumptions 4.1 and 4.2 hold. There exists an optimal stationary policy μ^* that always follows the LLLP principle. That is, at every system state $\mathbf{x} \in \mathcal{X}$, if the ith component of $\mu^*(\mathbf{x})$ is 1 (vehicle i is charged), then for every vehicle j such that $i \preccurlyeq j$ at $\mathbf{x}$, the jth component of $\mu^*(\mathbf{x})$ must also be 1 (vehicle j must be charged).

The proof of Theorem 4.2 is given in Appendix B. Since the state space is finite, there exists an optimal stationary policy (cf. page 175 of [2]). The crux of our proof centers on showing that any optimal stationary policy can be mapped to an optimal stationary policy that follows the LLLP principle, through an optimality condition (Bellman's equation) based argument.

V. Numerical Results

In this section, we compare the performance of three stationary heuristic policies, the EDF (Earliest Deadline First) policy and two LLF (Least Laxity First)-based heuristic policies. We consider a case with 400 chargers, i.e., $N = 400$ (large enough to accept all arriving vehicles in our simulation). The state of grid reflects the maximum capacity available for PHEV charging, i.e., the cost function associated with $s \in \mathcal{S}$ is given by

$$C(A, s) = 0, \quad \text{if } A \leq s; \qquad C(A, s) = Nq(E), \quad \text{if } A > s,$$

[8]This special state is recurrent in the Markov chain induced by every stationary policy [1], and therefore the minimum average cost with any initial state cannot be different from that with the special state being the initial state. Due to the space limit we omit the detailed proof and readers can refer to the proof of Prop. 7.4.1 in [1].

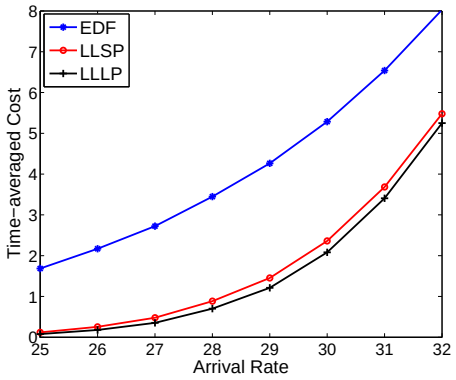

Fig. 1. A simulation experiment with $1,500,000$ trajectories for each arrival rate on the horizontal axis and time-averaged cost on the vertical axis, with non-completion penalty $q(n) = n$.

where $Nq(E)$ is an upper bound on the highest possible non-completion penalty that could incur to all vehicles in the set $\mathcal{I}_t$. Obviously, the operator should never charge more than s_t vehicles at stage t. The states of grid, $\{s_0, s_1, \ldots\}$, are assumed to be independent and identically distributed random variables that are uniformly distributed over $\mathcal{S} = \{40, 41, \ldots, 160\}$. Since we have assumed zero charging cost as an approximation for the case where the charging cost is much smaller than the non-completion penalty, in our simulation the only source of cost is non-completion penalty.

For simplicity, we consider a case where the number of arriving vehicles is a time-invariant constant, and the initial states of arriving vehicles are independent and identically distributed random variables. In particular, the number of stages for which a newly arrived vehicle i will stay at a charging facility, $\beta_i - \alpha_i$, is uniformly distributed over the set $\{1, \ldots, 10\}$ (i.e., $B = 10$), and the time needed to fulfill its request, γ_{i,α_i}, is uniformly distributed over the set $\{1, \ldots, \beta_i - \alpha_i\}$.

For a system state $\mathbf{x}_t$, let $V(\mathbf{x}_t)$ be the number of vehicles in the set $\mathcal{I}_t$ that are not fully charged. At a system state $\mathbf{x}_t$, the stationary EDF policy charges the first $\min\{s_t, V(\mathbf{x}_t)\}$ vehicles with the earliest departure times. For two vehicles that have the same deadline, π charges the one with less laxity. At a system state $\mathbf{x}_t$, both LLF-based policies charge the first $\min\{s_t, V(\mathbf{x}_t)\}$ vehicles with the least laxity. For two vehicles with the same laxity, the LLSP (Least Laxity and Shorter remaining Processing time) policy gives priority to the vehicle with shorter remaining processing time (an earlier departure time), while the LLLP (Least Laxity and Longer remaining Processing time) policy gives priority to the vehicle that has longer remaining processing time.

The time-averaged cost resulting from the three heuristic policies (EDF, LLSP, and LLLP) are compared in Fig. 1 and Fig. 2, for two different non-completion penalty functions $q(n) = n$ and $q(n) = n^2$. For both

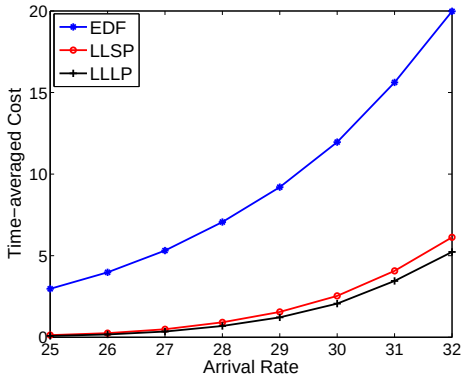

Fig. 2. A simulation experiment with $1,500,000$ trajectories for each arrival rate on the horizontal axis and time-averaged cost on the vertical axis, with non-completion penalty $q(n) = n^2$.

penalty functions, the numerical results show that the LLLP policy achieves the lowest time-averaged cost, and that LLSP significantly outperforms EDF. We note that the performance gap between LLSP and LLLP is much more significant under quadratic non-completion penalty. This is because the LLLP policy distributes the total remaining processing time to a larger number of vehicles with smaller remaining processing times, which in turn leads to lower non-completion penalty when the penalty function $q()$ is strictly convex. Indeed, under the linear penalty function, the LLLP policy reduces the time-averaged cost by $15\% - 35\%$ (compared to the LLSP policy) when the arrival rate is less than 30; while under the quadratic penalty function, the performance gap between LLSP and LLLP is much larger, and remains above 15% (of the cost resulting from LLSP) even when the arrival rate ranges in $\{30, 31, 32\}$.

VI. Conclusion

We formulate the scheduling problem of charging multiple PHEVs as a Markov decision process. Using an interchange argument, we prove the less laxity and longer remaining processing time (LLLP) principle: priority should be given to vehicles that have less laxity and longer remaining processing times, if the non-completion penalty function is convex and the operator does not discount future cost. We note that the LLLP principle is a partial characterization on the optimal scheduling policy, and that there may exist many stationary policies that do not violate the LLLP principle. A plausible future research direction is to compare and rank these heuristic policies in stylized models with more structures in system dynamics.

Appendix A
Proof of Lemma 4.1

To argue the feasibility of the interchanging policy $\bar{\pi}$, we will show that

1) in period w when the original policy π first charges vehicle j but not i, it is feasible for the interchanging policy to charge vehicle i;
2) in period $k = t+1, \ldots, w-1$, whenever the original policy π charges vehicle j, it is feasible for the interchanging policy to charge vehicle j.

It is straightforward to check the first point, i.e., at the stage w it is feasible for the interchanging policy $\bar{\pi}$ to charge vehicle i. This is because the original policy π charges vehicle i at stage t but the interchanging policy does not, and the interchanging policy does not charge vehicle i whenever the original policy π does not, at every stage before w.

We now prove the second point. We first consider the case where the original policy π first charges vehicle j but not i at stage w. Since at state $\mathbf{x}_t$, vehicle j has priority over i, and the policy π charges vehicle i but not j, we must have $\theta_{j,t+1} < \theta_{i,t+1}$ and $\gamma_{j,t+1} > \gamma_{i,t+1}$ at state $\mathbf{x}_{t+1}$. Before stage w (for $k = t+1, \ldots, w-1$), whenever the policy π charges vehicle j at state $\mathbf{x}_k$, it also charges vehicle i. It follows that for $k = t+1, \ldots, w-1$, whenever the policy π charges vehicle j at $\mathbf{x}_k$, we must have $\gamma_{j,k} > \gamma_{i,k} \geq 1$, which implies that $\hat{\gamma}_{j,k} = \gamma_{j,k} - 1 \geq 1$, i.e., it is feasible for the interchanging policy to charge j.

A similar argument applies to the case where the set $\mathcal{G}(\{\mathbf{x}_k\}_{k=t+1}^{t+W})$ is empty ($w = \infty$), and the deadline of vehicle j is no later than i's. We have $W = \beta_i - 1$. Before vehicle i's departure (for $k = t+1, \ldots, \beta_i - 1$), whenever the policy π charges vehicle j at state $\mathbf{x}_k$, it also charges vehicle i. It follows that for $k = t+1, \ldots, \beta_i - 1$, whenever the policy π charges vehicle j at $\mathbf{x}_k$, we have $\gamma_{j,k} > \gamma_{i,k} \geq 1$, and therefore $\hat{\gamma}_{j,k} = \gamma_{j,k} - 1 \geq 1$.

We finally consider the case where $w = \infty$ and the deadline of j is later than i's. We have $W = \beta_j - 1$. At state $\mathbf{x}_{\beta_i - 1}$, we know vehicle j's laxity must be strictly less than vehicle i's, i.e.,

$$\theta_{j,\beta_i-1} = \lambda_{j,\beta_i-1} - \gamma_{j,\beta_i-1} < \lambda_{i,\beta_i-1} - \gamma_{i,\beta_i-1} = 1 - \gamma_{i,\beta_i-1}.$$

If $\gamma_{i,\beta_i-1} \geq 1$, then we have $\theta_{j,\beta_i-1} < 0$. If $\gamma_{i,\beta_i-1} = 0$, since the policy π does not charge vehicle i at state $\mathbf{x}_{\beta_i-1}$, it does not charge vehicle j, because the set $\mathcal{G}(\{\mathbf{x}_k\}_{k=t+1}^{t+W})$ is empty. In this case, we have $\theta_{j,\beta_i-1} \leq 0$ and $\theta_{j,\beta_i} < 0$. In either case, for $k = \beta_i, \ldots, \beta_j - 1$, we have $\theta_{j,k} < 0$. It follows that for $k = \beta_i, \ldots, \beta_j - 1$, $\hat{\gamma}_{j,k} = \gamma_{j,k} - 1 \geq 1$.

APPENDIX B
PROOF OF THEOREM 4.2

We first introduce a necessary and sufficient condition for the optimality of a stationary policy. This condition will be used later in the proof. Under Assumption 4.2, the minimum average cost λ, together with an $|\mathcal{X}|$-dimensional vector $\mathbf{h} = \{h(\mathbf{x})\}_{\mathbf{x}\in\mathcal{X}}$, satisfies the following Bellman's equation:

$$\lambda + h(\mathbf{x}) = \min_{\mathbf{a}\in U(\mathbf{x})} \left\{ g(\mathbf{x},\mathbf{a}) + \sum_{\mathbf{y}\in\mathcal{X}} p_{\mathbf{x},\mathbf{y}}(\mathbf{a}) h(\mathbf{y}) \right\}, \quad \forall \mathbf{x} \in \mathcal{X}, \tag{7}$$

where $U(\mathbf{x})$ denotes the set of feasible actions at system state $\mathbf{x}$. It is known that a stationary policy μ^* is optimal, i.e., $J_{\mu^*}(\mathbf{x}) = \lambda$ for all $\mathbf{x}$, if $\mu^*(\mathbf{x})$ attains the minimum in (7) for each $\mathbf{x} \in \mathcal{X}$ [2].

Let $(\lambda, \mathbf{h})$ be a solution to the Bellman's equation in (7). For every $\mathbf{x} \in \mathcal{X}$, $h(\mathbf{x})$ (usually referred to as the differential cost for state $\mathbf{x}$) is the minimum, over all policies, of the difference between the expected cost to reach the special state $\bar{\mathbf{x}}$ (cf. (6)) from $\mathbf{x}$ for the first time and the cost that would be incurred if the cost per stage were equal to the minimum average λ at all states. To formally define $h(\mathbf{x})$, consider a modified dynamic program (DP) that is the same as the DP formulated in Section III, except that the special system state is absorbing ($\tilde{p}_{\bar{\mathbf{x}},\bar{\mathbf{x}}}(\mathbf{a}) = 1$ for every action vector $\mathbf{a}$), and the cost associated with the special system state is λ ($\tilde{g}(\bar{\mathbf{x}},\mathbf{a}) = \lambda$ for every feasible action $\mathbf{a}$). Here, $\tilde{p}$ and $\tilde{g}$ denote the state transition and stage cost functions of the modified Markov decision process, respectively. The differential cost for state $\mathbf{x}$ can be written as

$$h(\mathbf{x}) \stackrel{\Delta}{=} \min_{\pi} \limsup_{T\to\infty} \mathbb{E}\left\{ \sum_{t=0}^{T-1} (\tilde{g}(\mathbf{x}_t, \nu_t(\mathbf{x}_t)) - \lambda) \,\middle|\, \mathbf{x}_0 = \mathbf{x} \right\}, \tag{8}$$

where $\pi = \{\nu_0, \nu_1, \dots, \}$. In the proof of Lemma 4.2 we have shown that for any initial state $\mathbf{x}$, the special system state is visited with positive probability at least once within the first $(B + \max\{m, |\mathcal{D}|\})$ steps, which implies that $h(\mathbf{x})$ is finite for every $\mathbf{x} \in \mathcal{X}$.

We next employ a simple interchange argument to prove the existence of an optimal stationary policy that follows the LLLP principle. Let μ be an optimal stationary policy, and suppose that there exists a system state $\mathbf{x}_t$ such that the ith component of $\mu(\mathbf{x}_t)$ is 1, the jth component of $\mu(\mathbf{x}_t)$ is 0, and that $i \preccurlyeq j$. Since $\gamma_{i,t} \geq 1$ and $i \preccurlyeq j$, we must have $\gamma_{j,t} \geq 1$. We argue that another stationary policy $\bar{\mu}$, which agrees with the decisions made by the policy μ except that $\bar{\mu}$ charges vehicle j instead of i at the system state $\mathbf{x}_t$, must also be optimal.

For the state $\mathbf{x}_t$, the optimal stationary policy μ attains the minimum in (7) [2]. To argue that the stationary policy $\bar{\mu}$ is optimal, we only need to show that the stationary policy $\bar{\mu}$ also attains the minimum

in (7), i.e.,

$$\begin{aligned} &g(\mathbf{x}_t, \bar{\mu}(\mathbf{x}_t)) + \sum\nolimits_{\bar{\mathbf{x}}_{t+1} \in \mathcal{X}} p_{\mathbf{x}_t, \bar{\mathbf{x}}_{t+1}}(\bar{\mu}(\mathbf{x}_t)) h(\bar{\mathbf{x}}_{t+1}) \\ &\leq g(\mathbf{x}_t, \mu(\mathbf{x}_t)) + \sum\nolimits_{\mathbf{x}_{t+1} \in \mathcal{X}} p_{\mathbf{x}_t, \mathbf{x}_{t+1}}(\mu(\mathbf{x}_t)) h(\mathbf{x}_{t+1}). \end{aligned} \tag{9}$$

For every $\mathbf{x}_{t+1}$ such that $p_{\mathbf{x}_t, \mathbf{x}_{t+1}}(\mu(\mathbf{x}_t)) > 0$, there exists a corresponding system state $\bar{\mathbf{x}}_{t+1}$ that occurs with equal probability under the policy $\bar{\mu}$. The system state $\bar{\mathbf{x}}_{t+1}$ is the same as $\mathbf{x}_{t+1}$ except that $\bar{\gamma}_{i,t+1} = \gamma_{i,t+1} + 1$ and $\bar{\gamma}_{j,t+1} = \gamma_{j,t+1} - 1$. For every system state $\mathbf{x}_{t+1}$ that occurs with positive probability under the policy μ, we will show that

$$g(\mathbf{x}_t, \bar{\mu}(\mathbf{x}_t)) + h(\bar{\mathbf{x}}_{t+1}) \leq g(\mathbf{x}_t, \mu(\mathbf{x}_t)) + h(\mathbf{x}_{t+1}), \tag{10}$$

which implies the result in (9).

For every system state $\mathbf{x}_{t+1}$ that occurs with positive probability under the policy μ, since the state space is finite, there exists a policy $(\nu_{t+1}, \nu_{t+2}, \ldots)$ that attains the minimum on the right hand side of (8) for the state $\mathbf{x}_{t+1}$. For the policy $\pi = (\mu, \ldots, \mu, \nu_{t+1}, \nu_{t+2}, \ldots)$ (a non-stationary policy that agrees with the stationary policy μ from stage 0 through t), consider its interchanging policy $\bar{\pi} = (\mu, \ldots, \mu, \bar{\mu}, \bar{\nu}_{t+1}, \bar{\nu}_{t+2}, \ldots)$ with respect to vehicles i and j of state $\mathbf{x}_t$ (cf. Definition 4.2). In the proof of Theorem 4.1, we have shown that compared to the original policy π, the interchanging policy $\bar{\pi}$ cannot increase the total cost realized on any trajectory (from stage t through stage $t+W$)[9] that would occur with positive probability (cf. Eq. (4)). It follows that

$$\begin{aligned} &g(\mathbf{x}_t, \bar{\mu}(\mathbf{x}_t)) + g(\bar{\mathbf{x}}_{t+1}, \bar{\nu}_{t+1}(\mathbf{x}_{t+1})) + \mathbb{E}_{\bar{\pi}} \left\{ \sum\nolimits_{k=t+2}^{t+W} g(\bar{\mathbf{x}}_k, \bar{\nu}_k(\bar{\mathbf{x}}_k)) \right\} \\ &\leq g(\mathbf{x}_t, \mu(\mathbf{x}_t)) + g(\mathbf{x}_{t+1}, \nu_{t+1}(\mathbf{x}_{t+1})) + \mathbb{E}_{\pi} \left\{ \sum\nolimits_{k=t+2}^{t+W} g(\mathbf{x}_k, \nu_k(\mathbf{x}_k)) \right\}, \end{aligned} \tag{11}$$

where the expectations are over future system states from stage $t+2$ through $t+W$ induced by the policy $\bar{\pi}$ and π, respectively. For stages $k = t+1, \ldots, t+W$, since at least one vehicle (i or j) is at a charging facility, the special system state[10] is not reached, and therefore $\tilde{g}(\mathbf{x}_k, \mathbf{a}_k) = g(\mathbf{x}_k, \mathbf{a}_k)$ for any action $\mathbf{a}_k$.

[9] We let $W = \max\{\lambda_{i,t}, \lambda_{j,t}\} - 1$.

[10] The special system state is introduced in Eq. (6), and the cost function $\tilde{g}(\mathbf{x}_k, \mathbf{a}_k)$ is defined prior to Eq. (8).

After stage $t+W$, the two policies π and $\bar{\pi}$ result in the same expected cost. We therefore have

$$\begin{aligned}
&g(\mathbf{x}_t, \bar{\mu}(\mathbf{x}_t)) + h(\bar{\mathbf{x}}_{t+1}) + \lambda \\
&\leq g(\mathbf{x}_t, \bar{\mu}(\mathbf{x}_t)) + \tilde{g}(\bar{\mathbf{x}}_{t+1}, \bar{\nu}_{t+1}(\mathbf{x}_{t+1})) + \limsup_{T\to\infty} \mathbb{E}_{\bar{\pi}} \left\{ \sum\nolimits_{k=t+2}^{t+T} (\tilde{g}(\bar{\mathbf{x}}_k, \bar{\nu}_k(\bar{\mathbf{x}}_k)) - \lambda) \right\} \\
&\leq g(\mathbf{x}_t, \mu(\mathbf{x}_t)) + \tilde{g}(\mathbf{x}_{t+1}, \nu_{t+1}(\mathbf{x}_{t+1})) + \limsup_{T\to\infty} \mathbb{E}_{\pi} \left\{ \sum\nolimits_{k=t+2}^{t+T} (\tilde{g}(\mathbf{x}_k, \nu_k(\bar{\mathbf{x}}_k)) - \lambda) \right\} \\
&= g(\mathbf{x}_t, \mu(\mathbf{x}_t)) + h(\mathbf{x}_{t+1}) + \lambda,
\end{aligned} \tag{12}$$

where the first inequality follows from the definition of $h(\bar{\mathbf{x}}_{t+1})$ in (8), the second inequality follows from (11), and the last equality is true because the policy $\pi = (\mu, \ldots, \mu, \nu_{t+1}, \nu_{t+2}, \ldots)$ attains the minimum on the right hand side of (8) with an initial state $\mathbf{x}_{t+1}$.

We have shown the inequality in (10) holds for for every $\mathbf{x}_{t+1}$ such that $p_{\mathbf{x}_t,\mathbf{x}_{t+1}}(\mu(\mathbf{x}_t)) > 0$. It follows from (9) that the stationary policy $\bar{\mu}$ attains the minimum in (7), and is therefore optimal. Since the state space is finite, by repeating this interchange argument for finitely many times, we can construct an optimal stationary policy that follows the LLLP principle at all system states.

REFERENCES

[1] D. P. Bertsekas, *Dynamic Programming and Optimal Control*, 3rd Edition, vol. I. Athena Scientific: Nashua, New Hampshire, 2011.

[2] D. P. Bertsekas, *Dynamic Programming and Optimal Control*, 3rd Edition, vol. II. Athena Scientific: Nashua, New Hampshire, 2011.

[3] P. P. Bhattacharya and A. Ephremides, "Optimal scheduling with strict deadlines," *IEEE Trans. on Automatic Control*, vol. 34, no.7, pp. 721-728, 1989.

[4] M. C. Caramanis and J. M. Foster, "Coupling of day ahead and real-time power markets for energy and reserves incorporating local distribution network costs and congestion," in *Proc. of IEEE Allerton Annual Conference*, Illinois, USA, 2010.

[5] P. P. Bhattacharya, L. Tassiulas, and A. Ephremides, "Optimal scheduling with deadline constraints in tree networks," *IEEE Trans. on Automatic Control*, vol. 42, no. 12, pp. 1703-1705, 1997.

[6] O. Boucher, "Comparison of physically- and economically-based CO2-equivalences for methane," *Earth System Dynamics*, vol. 3, pp. 1-29, 2012.

[7] K. Clement-Nyns, E. Haesen, and J. Driesen, "The impact of charging plug-in hybrid electric vehicles on a residential distribution grid," *IEEE Trans. Power Systems*, vol. 25, no. 1, pp. 371-380, 2010.

[8] R. I. Davis, and A. Burns, "A survey of hard real-time scheduling for multiprocessor systems," *ACM Computing Surveys (CSUR)*, vol. 43, no. 4, 2011.

[9] M. Dertouzos, "Control robotics: the procedural control of physical processes," in *Proc. of IFIP Congress*, pp. 807-813, 1974.

[10] M. L. Dertouzos and A. K. Mok, "Multiprocessor online scheduling of hard-real-time tasks," *IEEE Transactions on Software Engineering,* vol. 5, pp. 1497-1506, 1989.

[11] L. Gan, U. Topcu, and S. Low, "Optimal decentralized protocols for electric vehicle charging," *IEEE Trans. Power Systems*, vol. 28, no. 2, pp. 940-951, 2013.

[12] C. Goebel and D. S. Callaway, "Using ICT-controlled plug-in electric vehicles to supply grid regulation in California at different renewable integration levels," *IEEE Trans. on Smart Grid*, vol. 4, no. 2, pp. 729-740, 2013.

[13] Q. Huang, Q. S. Jia, Z. Qiu, X. Guan, and G. Deconinck, "Matching EV Charging Load With Uncertain Wind: A Simulation-Based Policy Improvement Approach," *IEEE Trans. on Smart Grid*, vol. 6, no. 3, pp. 1425-1433, 2015.

[14] E. L. Karfopoulos, and N. D. Hatziargyriou, "A multi-agent system for controlled charging of a large population of electric vehicles," *IEEE Transactions on Power Systems*, vol. 28, no.2, pp. 1196-1204, 2013.

[15] C. L. Liu and J. W. Layland, "Scheduling algorithms for multiprogramming in a hard-real-time environment," *J. of ACM*, vol. 20, pp. 46-61, 1973.

[16] C. D. Locke, "Best-effort decision-making for real-time scheduling," *PhD thesis*, Carnegie Mellon University, 1986.

[17] J. Lopes, F. Soares, and P. Almeida, "Integration of electric vehicles in the electric power system," in *Proc. of IEEE*, vol. 99, no. 1, pp. 168-183, 2011.

[18] Z. Ma, D. S. Callaway, and I. Hiskens, "Decentralized charging control for large populations of plug-in vehicles," *IEEE Trans. Control System Technology*, vol. 28, no.1 , pp. 67-78, 2013.

[19] S. J. Moura, H. K. Fathy, D. S. Callaway, and J. L. Stein, "A stochastic optimal control approach for power management in plug-in hybrid electric vehicles," *IEEE Trans. Control System Technology*, vol. 19, no. 3, pp. 545-555, 2011.

[20] N. Rotering and M. D. Ilić, "Optimal charge control of plug-in hybrid electric vehicles in deregulated electricity markets," *IEEE Trans. Power Systems*, vol. 26, no. 3, pp. 1021-1029, 2008.

[21] A. Subramanian, M. Garcia, A. Domínguez-García, D. Callaway, K. Poolla, and P. Varaiya, "Real-time scheduling of deferrable electric loads," in *Proc. of 2012 American Control Conference*, Montreal, Canada, 2012.

[22] D. Wu, D. C. Aliprantis, and L. Ying, "Load scheduling and dispatch for aggregators of plug-in electric vehicles," *IEEE Trans. on Smart Grid*, vol. 3, no. 1, pp. 368-376, 2012.

[23] Y. Xu and F. Pan, "Scheduling for charging plug-in hybrid electric vehicles," in *Proc. of IEEE 51st Annual Conference on Decision and Control (CDC)*, 2012.